

Accelerating the convergence of replica exchange simulations using Gibbs sampling and adaptive temperature sets

Thomas Vogel* and Danny Perez†

Theoretical Division (T-1), Los Alamos National Laboratory, Los Alamos, NM 87545

We recently introduced a novel replica-exchange scheme in which an individual replica can sample from states encountered by other replicas at any previous time by way of a global configuration database, enabling the fast propagation of relevant states through the whole ensemble of replicas. This mechanism depends on the knowledge of global thermodynamic functions which are *measured* during the simulation and not coupled to the heat bath temperatures driving the individual simulations. Therefore, this setup also allows for a continuous adaptation of the temperature set. In this paper, we will review the new scheme and demonstrate its capability. The method is particularly useful for the fast and reliable estimation of the microcanonical temperature $T(U)$ or, equivalently, of the density of states $g(U)$ over a wide range of energies.

INTRODUCTION

Simulating physical systems exhibiting multifaceted structural transitions (encountered, for example, in many protein folding problems) or containing extremely hard to find but thermodynamically important states (e.g., crystalline states in hard materials) is a notorious challenge. Replica-exchange (RE) and multicanonical (MUCA) sampling schemes have both been applied, in conjunction with Monte Carlo (MC) and molecular dynamics (MD) methods, to address this issue. They enable the investigation of systems at multiple or varying temperatures, such facilitating the exploration of the configurational space. One of the main practical challenges in the case of RE is the determination of good temperature sets for the heat baths. To this day, many sophisticated schemes have been proposed to solve this problem, see, e.g., [1–5]. Another limitation of generic RE schemes is that exchange of configurations obviously does not, on its own, change the ensemble of configurations present in the simulation. This is not optimal if the simulation is not converged yet, e.g., if some replicas have yet to locate all thermodynamically relevant states. In this case, some sort of population control, where “good” configurations can be multiplied and “irrelevant” ones deleted, can help (see [6] for an example).

MUCA simulations offer a different strategy to efficiently sample configuration space. In the case of MD, MUCA sampling simply consists of rescaling the interatomic forces: $f_{\text{muca}} = (T_0/T(U))f_{\text{can}}$ [7, 8], where $f_{\text{can}} = -\nabla U$ are the conventional forces and $T(U)$ is the microcanonical temperature. The difficulty lies in the estimation of this *a priori* unknown function. Metadynamics [9, 10] and statistical temperature molecular dynamics (STMD) [11] were successfully applied to this problem. However, these methods can show dynamical artifacts like hysteresis when the system is driven over first-order like transitions, for example. Further, even with the knowledge of $T(U)$, the diffusion of a MUCA simulation in energy space can be very slow. However,

a very powerful side-effect of MUCA simulations is that, once obtained, $T(U)$ can be used to infer the density of states, hence providing a wealth of thermodynamic information.

In the following, we review a generalized RE method, the Gibbs-sampling enhanced RE, which addresses some of the issues with generic RE and MUCA simulations. By uncoupling the measured microcanonical temperature from the heat-bath temperatures which drive the simulations we can introduce powerful conformational mixing steps and continuously adapt the heat-bath temperatures for a complete coverage of the whole energy range. We demonstrate this ability by presenting a particular adaptive simulation which is set up so that only a tiny part of the phase space is covered the beginning. Note that $T(U)$ is also available at the end of such a simulation, providing an alternative to conventional MUCA approaches.

METHOD REVIEW

We consider a RE scheme [12] in an expanded-canonical ensemble [13] with reference temperatures T_0^i . In contrast to other sampling methods where thermodynamic functions are continuously adapted [8, 9, 14], the basic idea in our method [15] is to *measure* the microcanonical temperature while running canonical simulations, hence decoupling the measurements from the simulation-driving heat bath temperatures. We use an estimator for the microcanonical temperature based the microcanonical averages of time-derivatives of the product between the normalized energy gradient $\eta = \nabla U/(\nabla U \nabla U)$ and the particle momenta p :

$$g = \frac{dF(U)}{dU} = - \left\langle \frac{d}{dt} (\eta \cdot p) \right\rangle, \quad (1)$$

where F is the free energy, a relation which was derived in a more general form in the context of the adaptive biasing force (ABF) method [16]. The time derivatives are estimated in an additional, constant-energy MD time

step. As defined, g relates to the microcanonical temperature as:

$$T_{\text{ABF}}(U) = \frac{T_0}{1 - g}. \quad (2)$$

During the simulation, all replicas contribute measurements to the global calculation of $T_{\text{ABF}}(U)$, while they also collectively fill a global conformational database.[17] An instantaneous estimator for the (unnormalized) entropy $S(U)$ can then be obtained:

$$S(U) = \int_{U_0}^U T_{\text{ABF}}^{-1}(U') dU' + C. \quad (3)$$

With $g(U) = \exp[k_B^{-1}S(U)]$ being the density of states, canonical energy probability distributions $P_i(U) \propto g(U)e^{-\beta_i U}$ can be calculated. The key is that these distributions calculated from global data are different from those that would be inferred from the data gathered by individual walkers alone. For example, $P_i(U)$ can account for the presence of some low energy state, even though walker i might never have sampled it himself. This allows for the introduction of a rejection-free, global Monte Carlo move facilitating the propagation of important states through the simulation: at any time, a walker can sample a new energy value from the distribution $P_i(U)$ (a so-called ‘‘Gibbs sampling’’ move) and replace its actual configuration by one with that energy ($\pm\epsilon$, for practical reasons) from the global configuration database. Such moves allow each replica to properly sample from the pool of states visited by all other replicas.

THE ADAPTIVE TEMPERATURE SET

In RE, good performance relies on a proper choice of temperatures, as this determines the exchange probabilities between neighboring replicas, say i and $i + 1$. Assuming the corresponding canonical distributions $P_i(U)$ and $P_{i+1}(U)$ are known, the exchange probabilities $W_i(P_i, P_{i+1})$ read:

$$W_i(P_i, P_{i+1}) = \int_{-\infty}^{\infty} P_i(U) \int_{-\infty}^{\infty} W_{\text{acc}}(U, U') P_{i+1}(U') dU dU', \quad (4)$$

with

$$W_{\text{acc}}(U, U') = \min \{1, \exp [(\beta_{i+1} - \beta_i)(U' - U)]\} \quad (5)$$

being the RE acceptance probability. These $W_i(P_i, P_{i+1})$ can be approximated (see [18], for example), but with the knowledge of a global, instantaneous $T(U)$ and hence of $g(U)$ and $P_i(U)$, the $W_i(P_i, P_{i+1})$ can be calculated. We can then find a temperature set $\beta_i^{-1} = k_B T_0^i$, ($i = 1, \dots, n$), such that [cf. [18, 19]]

$$W_i(P_i, P_{i+1}) = \text{const} \quad \forall i. \quad (6)$$

Such a criterion would be optimal in the limit where sampling is ergodic on the timescale of exchange attempts [18, 19]. When this condition is broken, it might instead be preferable to directly minimize the tunneling time [2].

Since there is no need for the reference temperatures T_0^i to remain constant for the calculation of $T_{\text{ABF}}(U)$ or other thermodynamic averages, these can be freely changed during the simulation. With a given, fixed number of processors, there are two options: either require a fixed value for the constant in Eq. 6 (resulting in a fluctuating total temperature range) or fix the global temperature range (and thus correspondingly adjusts the value of the constant). We chose the latter to ensure we always cover the whole temperature range and adjust the temperatures using a bisection scheme.

DEMONSTRATION

We illustrate the performance of the method for a system of 500 silver atoms interacting via an embedded atom potential [20]. We use a constant particle density of $\rho = 0.0585 \text{ \AA}^{-3}$ and perform a Gibbs-sampling enhanced RE simulation based on molecular dynamics runs at heat-bath temperatures T_0^i . Representative configurations are shown in Fig. 1. We begin the simulation in an extremely unnatural way: of the 71 individual canonical walkers, half of them start at the minimum temperature of $T_0^{i < 35} = T_{\text{min}} = 100 \text{ K}$ and the other half at the maximum temperature $T_0^{i \geq 35} = T_{\text{max}} = 3500 \text{ K}$. Furthermore, all replicas begin from an amorphous configuration. Figure 2 shows snapshots of the instantaneous microcanonical temperature measured as $T_{\text{ABF}}(U)$ (cf. Eq. 2) for different simulation times t_{MD} . Shortly after starting the simulation (Fig. 2a) and before the first temperature adaptation, we only find configurations with temperatures scattered around the initial values T_{min} and T_{max} . The instantaneous temperature $T_{\text{ABF}}(U)$ is completely arbitrary at this point.[21] The next snapshot (b) was taken after the first temperature adaptation and the exploration of configurations at intermediate energies starts. After about 5×10^5 MD steps (Fig. 2c;




FIG. 1. Representative configurations of the 500-silver atom system. From left to right: an amorphous state; the pure crystalline (ground-) state; two crystalline states with stacking faults.

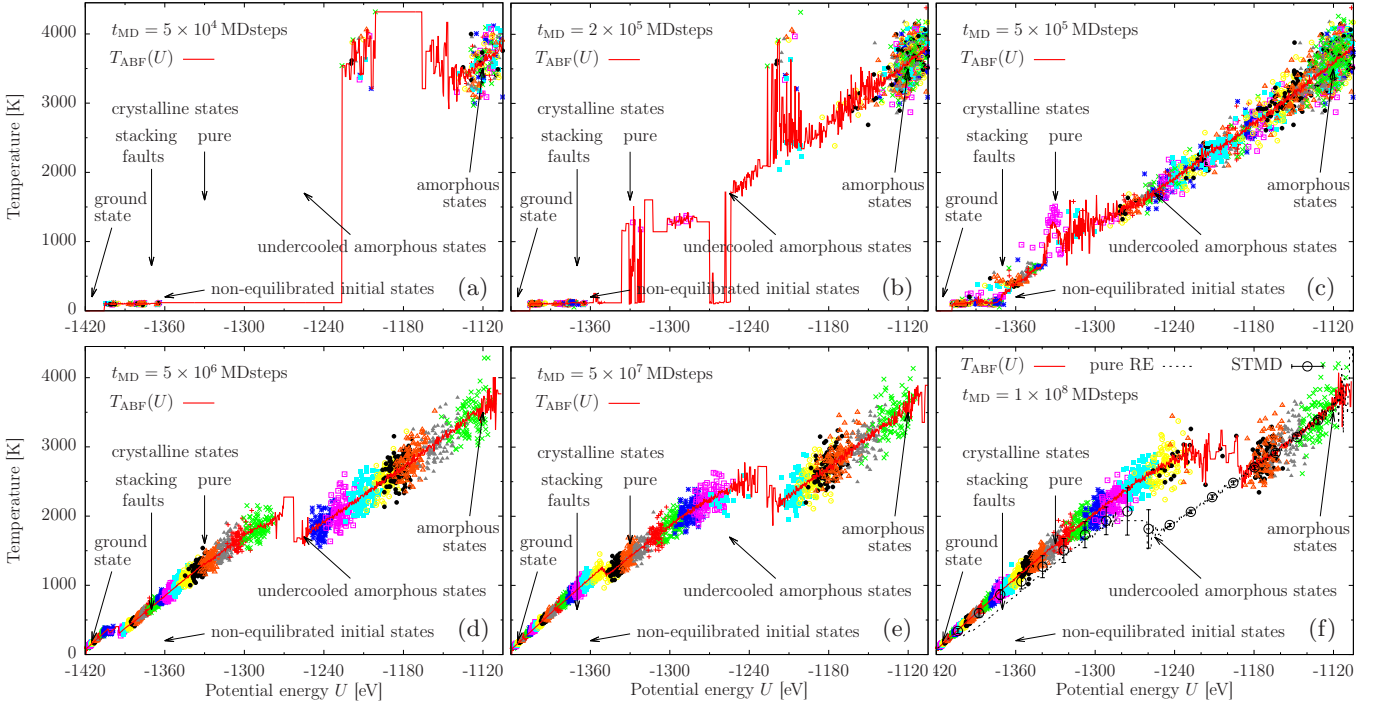


FIG. 2. Snapshots of the instantaneous estimator for the microcanonical temperature (red, solid line) at different times during the simulation. The dots correspond to raw measurements entering in the average (cf. Eqs. 1 and 2). Different colors and symbols correspond to different reference temperatures T_0^i . Raw measurements are only shown for every other reference temperature. Only 10% of the raw data are shown for clarity. The last plot (f) also contains data from pure RE and from STMD (see text for details).

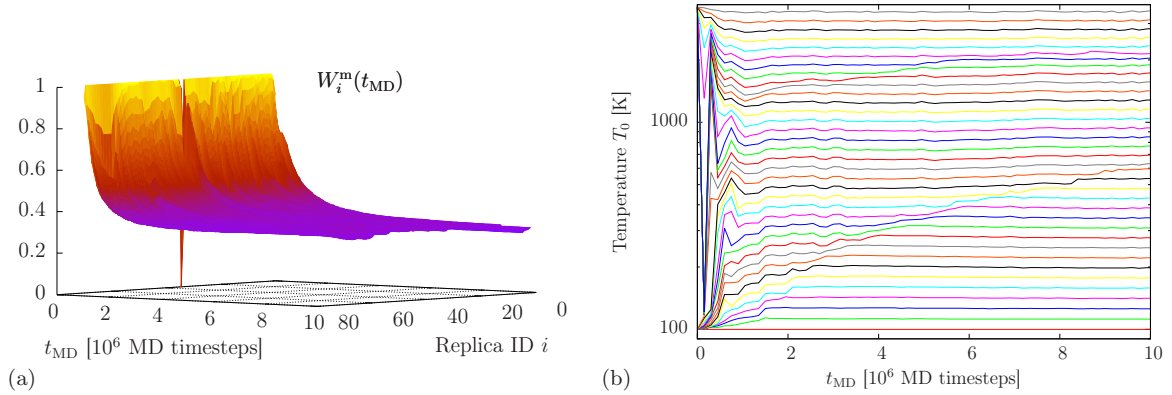


FIG. 3. (a) Measured replica-exchange acceptance rates for all neighbor pairs $(i, i + 1)$ over time. (b) Time evolution of the temperature set T_0^i .

temperature adaptation takes place every 1.5×10^5 MD steps) temperatures are already spread out such that the whole energy range is covered. The ground state has been discovered after 5×10^6 MD steps (d) and propagates through the simulation (e) facilitated by the Gibbs sampling moves. The simulation eventually converges after approximately 1×10^8 MD steps (f). Note that we set a memory time after which data is discarded, so that non-equilibrated states encountered in the beginning do not enter in the averages (Eqs. 1 and 2) at later times.

We show in Fig. 3 the *measured* replica-exchange acceptance rate W_i^m (a) and the values of the heat bath temperatures T_0^i (b) as the simulation time progresses. In accordance to the discussion above, $W_i^m = 0$ for the exchange between the two neighboring walkers at T_{\min} and T_{\max} and $W_i^m = 1$ everywhere else. After a short time, however, all measured acceptance rates approach a constant value of $W_i^m \approx 0.35$. This value is in agreement with the calculated value for the constraint given in Eq. 6 for our particular set-up, i.e., for the fixed temper-

ature range and the number of walkers we employ. Note that we can directly control the value of W_i^m by deploying more or less walkers. The adaptation speed of the replica-exchange rates is of course related to the adaption speed of the temperature set T_0^i . It can be seen in Fig. 3 b that very few temperature adaptation steps are required before they are reasonably well distributed. Adaptation is marginal at times $t_{MD} \gtrsim 1 \times 10^6$ and is mainly in response to the propagation of the ground-state, as this creates an artificial signal of a “phase transition” in $T_{ABF}(U)$ and the temperatures get locally closer together around the corresponding “transition temperature”.

Finally, in Fig. 2 f, we also provide data from a pure RE exchange run (i.e., w/o Gibbs moves) at the same time and from multiple Gauss-kernel STMD [8] runs, for comparison. While the pure RE run did not converge at all after this time and mainly samples configurations with stacking faults in the crystalline phase, the STMD run samples both, pure and defected states but suffers from hysteresis[22]. While both conventional methods are not able to converge in this challenging situation, our method including additional Gibbs-sampling moves works remarkably well.

We acknowledge funding by Los Alamos National Laboratory’s (LANL) Laboratory Directed Research and Development ER program. LANL is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. DOE under Contract DE-AC52-06NA25396. Assigned LA-UR-15-21704.

* tvogel@lanl.gov; tvogel1@stetson.edu

† danny_perez@lanl.gov

- [1] N. Rathore, M. Chopra, and J. J. de Pablo, J. Chem. Phys. **122**, 024111 (2005).
- [2] H. G. Katzgraber, S. Trebst, D. A. Huse, and M. Troyer, J. Stat. Mech.: Theory Exp. **2006**, P03018 (2006).
- [3] A. Patriksson and D. van der Spoel, Phys. Chem. Chem. Phys. **10**, 2073 (2008).
- [4] M. Guidetti, V. Rolando, and R. Tripiccion, J. Comput. Phys. **231**, 1524 (2012).
- [5] A. J. Ballard and D. J. Wales, J. Chem. Theory Comput. (JCTC) **10**, 5599 (2014).
- [6] H.-P. Hsu and P. Grassberger, J. Stat. Phys. **144**, 597 (2011).
- [7] U. H. E. Hansmann, Y. Okamoto, and F. Eisenmenger, Chem. Phys. Lett. **259**, 321 (1996).
- [8] C. Junghans, D. Perez, and T. Vogel, J. Chem. Theory Comput. (JCTC) **10**, 1843 (2014).
- [9] A. Laio and M. Parrinello, Proc. Natl. Acad. Sci. **99**, 12562 (2002).
- [10] J. F. Dama, M. Parrinello, and G. A. Voth, Phys. Rev. Lett. **112**, 240602 (2014).
- [11] J. Kim, J. E. Straub, and T. Keyes, Phys. Rev. Lett. **97**, 050601 (2006).
- [12] C. J. Geyer, in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, edited by E. M. Keramidas (Interface Foundation, Fairfax Station, VA, 1991) p. 156.
- [13] A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992).
- [14] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
- [15] T. Vogel and D. Perez, Phys. Rev. Lett. **115**, 190602 (2015).
- [16] E. Darve, D. Rodríguez-Gómez, and A. Pohorille, J. Chem. Phys. **128**, 144120 (2008).
- [17] The management of all global data is done by an additional master process which does not perform a MD run himself.
- [18] W. Nadler and U. H. E. Hansmann, Phys. Rev. E **75**, 026109 (2007).
- [19] E. Bittner, A. Nußbaumer, and W. Janke, Phys. Rev. Lett. **101**, 130603 (2008).
- [20] P. L. Williams, Y. Mishin, and J. C. Hamilton, Model. Simul. Mater. Sci. Eng. **14**, 817 (2006).
- [21] If there are no measured values in particular energy bins, we simply copy the closest valid value from a lower energy.
- [22] This behavior is known and simply illustrates the perils of choosing a reaction variable that involves first-order-like transitions.